



Rec'd PCT/PTO 21 DEC 2001

The
Patent
Office

G-BOOK
08303

PCT/GB 00/02303

INVESTOR IN PEOPLE



HJ

10/019172

**PRIORITY
DOCUMENT**

SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)

The Patent Office
Concept House
Cardiff Road

Newport REC'D 18 AUG 2000
South Wales
NP10 8QQ WIPO

PCT

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated: 31 July 2000

THIS PAGE BLANK (USPTO)

THE PATENT OFFICE
D
26 JUN 1999
RECEIVED BY POST

The
**Patent
Office**

26 JUN 1999 E457456-1 D02982
P01/7700 0.00 - 9914876.9

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

The Patent Office

Cardiff Road
Newport
Gwent NP9 1RH

1. Your reference

DWS/VV/P238GB

2. Patent application number

(The Patent Office will fill in this part)

9914876.9

3. Full name, address and postcode of the or of each applicant (*underline all surnames*)

UNIVERSITY OF YORK
HESLINGTON
YORK
YO1 5DD.

Patents ADP number (*if you know it*)

0 597854 9001

If the applicant is a corporate body, give the country/state of its incorporation

4. Title of the invention

DATA PROCESSORS

5. Name of your agent (*if you have one*)

"Address for service" in the United Kingdom to which all correspondence should be sent
(including the postcode)

DAVID STANLEY INTELLECTUAL PROPERTY
STANLEY'S INTELLECTUAL PROPERTY
Kings Court
12 King Street
LEEDS
LS1 2HL

Patents ADP number (*if you know it*)

06871289002 ✓

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (*if you know it*) the or each application number

Country

Priority application number
(if you know it)

Date of filing
(day / month / year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (*Answer 'Yes' if:*
a) any applicant named in part 3 is not an inventor, or
b) there is an inventor who is not named as an applicant, or
c) any named applicant is a corporate body.
See note (d))

YES

THIS PAGE BLANK (USPTO)

9. Enter the number of sheets for any of the following items you are filing with this form.
Do not count copies of the same document

Continuation sheets of this form

Description	32
Claim(s)	6
Abstract	
Drawing(s)	5 <i>as</i>

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (*Patents Form 7/77*)

Request for preliminary examination and search (*Patents Form 9/77*)

1

Request for substantive examination
(*Patents Form 10/77*)

1

Any other documents
(please specify)

11.

I/We request the grant of a patent on the basis of this application.

Signature

Date

DAVID STANLEY INTELLECTUAL PROPERTY

25 June 1999

12. Name and daytime telephone number of person to contact in the United Kingdom

David Stanley

01481 824411

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- a) If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- b) Write your answers in capital letters using black ink or you may type them.
- c) If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- d) If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- e) Once you have filled in the form you must remember to sign and date it.
- f) For details of the fee and ways to pay please contact the Patent Office.

DUPLICATE

- 1 -

DATA PROCESSORS

This invention relates to data processors, and is concerned particularly although not exclusively with the matching of and subsequent
5 processing of data.

In known data matching processors, one enters query data to be matched with existing known data. To find a match, the query data is compared with all known existing data until a match is found. This can be a slow process, even with substantial processors.

10 Thus, a problem with such processors is in achieving data matching and processing at an acceptable speed, with realistic resources. Many known processors organise stored data in fields (as in a database), so that in order to find a stored record, the data in the query must also be organised into fields (e.g. street name, town, postcode, etc.). In many cases, the field
15 names may be unknown. For example, in entering as query data a postal address to be matched, one may not know if a word belongs to a "street" field or a "town" field.

Preferred embodiments of the invention aim to provide data processors that provide rapid matching of query data that can be of much
20 greater length, and which can remove the need for additional identifiers such as field names. In other words, query data can be entered, largely irrespective of order.

According to one aspect of the present invention, there is provided a data processor comprising:

25 a) a correlation matrix memory arranged to store data;

- b) input means arranged to receive sets of input data to be stored in the correlation matrix memory;
- c) a sampler arranged to derive, from each set of input data, a respective set of tuples;
- 5 d) a coder arranged to code each of the tuples;
- e) a combiner arranged to combine the coded tuples for a respective set of input data;
- f) a separator generator arranged to generate for each set of input data a respective, associated, unique separator;
- 10 g) storage means arranged to store the association of each separator with its respective set of input data; and
- h) addressing means arranged to applying to the correlation matrix memory, for each set of input data, the respective combined coded tuples as a row address and the respective unique separator as a column address, or vice-versa.

15 Preferably, the combined coded tuples for each set of input data are in the form of a binary coded word; the data processor further comprises a translator arranged to translate each such binary coded word into a translated word comprising index values representing which bits of the binary coded word are set; and said addressing means is arranged to apply the translated word to the correlation matrix memory.

20 Preferably, said separator generator is arranged to generate separators in a random manner.

Preferably, said separator generator is arranged to generate separators which are M bits wide and having N bits set, where N>1 or N=1, and where N<M.

Preferably, for each said set of tuples, each tuple comprises three
5 successive elements of a respective set of input data, and each successive tuple is offset by one such element from the preceding tuple.

Preferably, said coder is arranged to code said tuples by tensoring.

Preferably, said combiner is arranged to combine the coded tuples for a respective set of input data, by superimposition.

10 Preferably, at least some of the rows (or columns) of the correlation matrix memory are represented by binary words, each of which represents the positions of each bit in the respective row (or column) which is set.

15 Preferably, said correlation matrix memory comprises a plurality of sub-correlation matrix memories; said addressing means is arranged to access a first one of said sub-correlation matrix memories and apply the combined coded tuples of a respective set of input data to that sub-correlation matrix memory unless a respective row (or column) of that sub-correlation matrix memory will become saturated by application of those tuples; and in the event of such prospective saturation, access successive
20 ones of the sub-correlation matrix memories until those tuples can be applied to a respective one of the sub-correlation matrix memories without such saturation.

A data processor according to any of the preceding aspects of the invention may be arranged to receive sets of query data to be matched with
25 sets of input data stored in the correlation matrix memory, and to derive, for

each set of query data, a respective set of coded tuples analogous to those derived for the original input data, and to apply to the correlation matrix memory, for each set of query data, the respective combined coded tuples as a row (or column) address: the data processor further comprising:

- 5 i) output means for outputting a raw superimposed separator which represents, for a respective set of query data, the number of rows (or columns) having a bit set by the applied combined coded tuples in each column (or row) represented by the raw superimposed separator;
- 10 j) threshold means arranged to convert the raw superimposed separator into a binary superimposed separator; and
- k) an extractor arranged to extract individual separators from the binary superimposed separator.

Preferably, said thresholding means sets an absolute threshold value, 15 and provides said binary superimposed separator as a word in which bits represent respective columns (or rows) of the correlation matrix memory, and each of those bits is set if the number of rows (or columns) having a bit set by the applied combined coded tuples in the respective column (or row) equals or is greater than said absolute threshold value.

20 Said thresholding means may determine a value k, and provide said binary superimposed separator as a word in which bits represent respective columns (or rows) of the correlation matrix memory, and are set for the k respective columns (or rows) having the highest number of rows (or columns) which have a bit set by the applied combined coded tuples in the 25 correlation matrix memory.

A data processor as above may further comprise back-checking means arranged to compare sets of recalled data, identified by respective separators extracted by said extractor, with original query data, in order to identify the set or sets of recalled data which matches best the original query
5 data.

A data processor according to any of the preceding aspects of the invention may be arranged to process sets of input data and query data in the form of postal addresses.

According to another aspect of the present invention, there is
10 provided a method of processing data, comprising the steps of:

- a) receiving sets of input data to be stored in a correlation matrix memory;
- b) deriving, from each set of input data, a respective set of tuples;
- c) coding each of the tuples;
- 15 d) combining the coded tuples for a respective set of input data;
- e) generating for each set of input data a respective, associated, unique separator;
- f) storing the association of each separator with its respective set of input data; and
- 20 g) applying to the correlation matrix memory, for each set of input data, the respective combined coded tuples as a row address and the respective unique separator as a column address, or vice-versa.

According to another aspect of the invention, there is provided a method of processing data comprising the steps of:

- a) receiving sets of query data to be matched with sets of input data stored in a correlation matrix memory by a method according to the preceding aspect of the invention;
- 5 b) deriving, for each set of query data, a respective set of coded tuples analogous to those derived for the original input data;
- c) applying to the correlation matrix memory, for each set of query data, the respective combined coded tuples as a row (or column) address;
- 10 d) outputting a raw superimposed separator which represents, for a respective set of query data, the number of rows (or columns) having a bit set by the applied combined coded tuples in each column (or row) represented by the raw superimposed separator;
- e) converting the raw superimposed separator into a binary superimposed separator;
- 15 f) extracting one or more individual separator from the binary superimposed separator; and
- g) identifying the or each respective original set of input data from association with the or each individual separator extracted from the binary superimposed separator.

Any of the above methods may be carried out by a data processor according to any of the preceding aspects of the invention.

Any of the above methods may incorporate any of the features disclosed in this specification, claims and/or drawings.

For a better understanding of the invention, and to show how embodiments of the same may be carried into effect, reference will now be made, by way of example, to the accompanying diagrammatic drawings, in which:

Figure 1 is an illustration of a correlation matrix memory (referred to herein for convenience as a "CMM");

Figure 2 illustrates tupling, tensoring and coding of a simple input word;

Figure 3 is an illustration similar to that of Figure 1, but showing an output array of summed values;

Figure 4 is a block diagram of a data processor comprising one example of an embodiment of the invention;

Figure 5 illustrates an example of steps carried out by the data processor of Figure 4, to locate data stored in the correlation matrix memory of the data processor;

Figure 6a shows a conventional implementation of a CMM line;

Figure 6b shows an alternative implementation of a CMM line;

Figure 7 shows line usage thresholds for compressed CMMs;

Figure 8 shows memory requirement for various sizes of CMM;

Figure 9 shows relative memory savings for various sized CMMs, using compression; and

Figure 10 is an illustration of an MBI ("Middle-Bit-Indexing") extractor.

The basic principle of operation of a correlation matrix memory (CMM) is illustrated in Figure 1. The CMM is illustrated as having a plurality of rows and columns. The intersection of each row and column in Figure 1 represents one bit of memory so that, conceptually, the CMM comprises many one-bit cells, each initially set to '0'. During a teaching step, during which data is stored in the CMM, every cell at the intersection of an active row and an active column is set to '1', regardless of its current state. Other cells are left unchanged.

Thus, in Figure 1, the intersections with a solid dot ● have previously been set to '1', whilst additional intersections with a hollow dot ○ are newly set by the illustrated row address pattern 001000010 and column address pattern 010000100.

To recall data from the CMM of Figure 1, a row address pattern is applied, and a row is addressed only when a '1' appears in the row address pattern. For each column, a sum is generated of all intersections set at '1' in all rows addressed by the applied row address pattern, and the sum is subsequently processed.

For example, if the row address pattern illustrated in Figure 1 is applied again as query data to the CMM, in a recall step, then the column sum pattern 020100210 as illustrated would be obtained, as representing the number of intersections set at '1' in each of the columns, for all of the addressed rows.

Figure 1 is a very simple example to illustrate the mode of operation. In

practice, one would expect the numbers of intersections set at '1' to be rather higher, for at least some of the columns.

An understanding of the configuration and use of a CMM such as illustrated in Figure 1 in an embodiment of the invention may be made easier 5 by use of a specific example. Thus, there will now be described an example of storing and retrieving a list of addresses in the CMM. It is to be understood that this is by way of example only, and data of any other type may be stored and retrieved in an analogous manner.

The address data used in this example, which has been the subject of 10 our experimental research, is in a data file containing postal addresses.

The first step is to pre-process the address data file to combine some redundant elements representing small variations in basically the same address, and to remove exact duplications so that each record is unique. The main objective here is to reduce the size of file needed to represent the data, 15 so that smaller CMMs can be used. The current implementation of pre-processing reduces the number of records that must be stored from 26 million in the address data file to just over 4.3 million records.

Even before pre-processing, the address data file already contains some groups of multiple addresses in abbreviated form, stored in a single 20 record. For example, the even or odd-numbered houses in a street may appear as "N-M Argyle Street" (say) in a single record. The procedures described here additionally allow consecutive numbers to be combined into an address range in a single record as well.

Note that multiple passes are normally required, because the 25 procedures are applied recursively. However, several records can be

combined in a single pass so that, for example, the eight records covering house numbers 21, 22, 23, 24, 25, 26, 27 and 28 can be combined to produce the new record "21.28 High Street ... " (for example), if the original records occur successively in the address database.

5 The following procedures are applied in pre-processing the address data file to obtain an input database file.

If N records are identical, (N-1) of those records are removed.

10 If a consecutive range of two or more records differs in only one word, and that word is not the postcode, then the records are replaced by a new combined record using the following syntax:

1. Numeric strings (natural numbers) whose value differs by exactly one across a consecutive sequence of (otherwise identical) records are replaced by a new combined string separated by the "." character. Thus, for example, "21 High Street, Jollytown JO5 2TR", "22 High Street, Jollytown JO5 2TR" and "23 High Street, Jollytown JO5 2TR" are replaced by the single record: "21.23 High Street, Jollytown JO5 2TR".
2. Numeric strings (natural numbers) whose value differs by exactly two across a consecutive sequence of (otherwise identical) records are replaced by a new combined string separated by the "-" character. Thus, for example, "21 High Street, Jollytown JO5 2TR", "23 High Street, Jollytown JO5 2TR" and "25 High Street, Jollytown JO5 2TR" are replaced by the single record: "21-25 High Street, Jollytown JO5 2TR"

3. Alphanumeric strings (but not postcode strings) are replaced by a new combined ("folded") record but separated by whitespace, and delimited by ">" and "<" characters. Thus, for example, "Old Farm Cottage, Ditherington, Blistershire BL6 4DI" and "New Farm Cottage, Ditherington, Blistershire BL6 4DI" are replaced by: "> Old New < Farm Cottage, Ditherington, Blistershire BL6 4DI".

5

All of the above procedures result in compressing the input data.

Lastly, records are transferred to the input database file in a pseudo-random order. The intended purpose of this re-ordering is to reduce the occurrence of clusters of similar text strings, by distributing these more uniformly through the input database. That is, the original address database is supplied in what is called 'Postcode Area Order', where the file is sorted according to the postcode. Addresses which share the same postcode are 15 then sorted according to other fields in the address database such as street name, building name and locality. This means that, for example, the first 3000 or so addresses in the address data file all belong to the AB10 postcode area - somewhere in Aberdeen. All of these records will therefore have a much higher degree of similarity than 3000 addresses taken at 20 random from the database. By taking the records in random order, there will tend to be a much wider variation of data presented to any particular CMM before it begins to get saturated. Saturation of CMMs will be discussed again below.

In order to enter data into a CMM for storage, the text of the 25 addresses has to be converted into binary code. In order to do this, each character is assigned a unique binary code. Figure 2(a) shows an example

of this, where six letters are each assigned a respective six-bit binary code, in each of which only a single bit is set to '1'.

The text to be stored is then subjected to n-tuple sampling. By this it is meant that each string of characters is divided into a succession of samples of n characters, each sample being one character on from the previous sample. Another way to look at this is as a "sliding window" n characters wide, which moves across a stream of input characters, such that the "window" advances one character at a time.

This is most readily understood by reference to Figure 2(b), where the word SPOTTER is sampled (for simplicity) as six 2-tuples, each comprising a successive pair of characters.

A unique binary code is then assigned to each 2-tuple or pair of characters, as the result of combining the binary codes of the individual characters of the pair, using a binary tensor product operation. An example of this is shown in Figure 2(c), where the tensor product of the first 2-tuple of Figure 2(b) is formed as the cross product of the two binary values of the letters 'S' and 'P'. The resulting matrix is then converted back into a single binary number simply by taking each row of the matrix in turn and concatenating them together. The order in which the rows are taken is not important, provided that it is consistent for all tensor products. It is important to note that the final single binary number resulting from the tensoring of each 2-tuple depends upon the order of the letters in the tuple – i.e. the result from "PS" is different to that from "SP".

When all of the 2-tuples in the text have been sampled and tensored to produce respective binary numbers (six in this simple example), all of those binary numbers are combined by a binary OR operation to produce a final

binary pattern or number for the text. Figure 2(d) illustrates the final binary pattern or number for the input word "SPOTTER". Figure 2(e) shows the corresponding final binary pattern or number for the slightly mis-spelt input word "SPOTER". Note the similarity between the binary patterns of Figures 5 2(d) and 2(e).

By contrast, Figure 2(f) shows the final binary pattern or number for the input word "PROTEST". Note the marked difference between the binary patterns of Figures 2(d) and 2(f), even though the word PROTEST is an anagram of the word SPOTTER, and therefore comprises exactly the same 10 characters.

Thus, the final patterns for two words which are anagrams of each other will usually be quite different, whereas a spelling mistake will only affect two tensor patterns (for 2-tuples), so the overall pattern for the word should not be overly disrupted.

15 Although 2-tupling and tensoring has been illustrated for simplicity, higher order tupling and tensoring can equally well be carried out. We have found 3-tupling and tensoring to be particularly efficient.

The illustrative example of Figure 2 shows tupling and tensoring of a single word. In the practical example of the address data file, each record 20 (i.e. a complete postal address) is sampled (typically in 3-tuples or "triples") and tensored to give a resultant binary number or pattern for that record. There is then generated for that record (typically using a pseudo-random process) a second binary pattern which we refer to as a separator, which is associated with the binary pattern of the record. The purpose of the 25 separator is to represent the record which may later be found by a search

operation, and so a record is maintained of the original data record represented by each separator.

The binary pattern for each record is then entered into the CMM with its respective separator. That is, the binary pattern for the record is applied 5 as a row address and its respective separator as a column address (or vice-versa), and the intersections of the CMM which have both row and column addressed as '1' are set as '1'. This is generally in the manner as described above with reference to Figure 1.

All desired records are entered into the CMM in this way, which may 10 be referred to as a step of "teaching" the CMM.

Once all of the desired data has been stored in the CMM, records can be recalled from the CMM when desired, by a "recall" or "search" step, as follows.

Firstly it is necessary to enter query data – that is, as much of an 15 original record as is available, to identify the full record. The query data is then processed in the same way as original data was entered in the CMM in the first place – that is, it is sampled in tuples, the tuples are binary coded and tensored, and the tensored products combined to form a final binary pattern of the query data.

20 The final binary pattern thus formed is then applied as a row address to the CMM. Then, for each column of the CMM, for each row of the address set to '1', the number of intersections of that row and column which are set to '1' are counted, to give a sum for that column. The sequence of sums for all of the columns gives a 1-dimensional output array of summed 25 values. As a very simple example of this, Figure 3 shows the CMM of

Figure 1, with the same row address as in Figure 1 applied, and showing the array of summed values of the columns as 020100210. This represents a number of combined separators that may match the input query represented by the row address applied to the CMM.

- 5 To reduce the number of potential "hits" represented by the separators included in the output array of summed values, a thresholding step is applied. For example, referring to Figure 2, if a threshold value of "2" is applied to the output array of summed values, each bit of the array which is equal to or greater than "2" is set to '1', and all other bits are set to '0'.
- 10 This then reduces the original output array of summed values to a binary output array. In the simple example of Figure 3, the binary output array is 01000010. Due to the simplicity of this example, this is identical to the original separator code applied in Figure 1, which then links to the original record. However, as a general rule, the binary output array will be more
- 15 complicated, representing a number of superimposed separators, and will require further processing by an extractor to extract the correct separator for the query data and therefore link to the original record that is sought.

- 20 It is quite likely that, given incomplete query data, there will be a number of possible matches to the query. In this case, the extractor will extract a number of individual separators, which are then linked to their respective records, which in turn can be listed, preferably in a ranking order.

The above description outlines a data processor as one example of an embodiment of the invention. The essential parts of such a data processor 1 are illustrated diagrammatically in Figure 4.

- 25 In Figure 4, query data is input at 10, having been subject to pre-processing as outlined above. The data is then passed to a sampler 2,

which forms n-tuples. The tuples are then coded. To this end, a tensoring means 3 forms tensor patterns from the binary codes of the letters of the tuples, and the tensor patterns are combined by a superimposing coder 4. The output of the coder 4 is translated by an index-value coder 5, which 5 outputs a series of values representing the positions of each respective bit which is set to '1' in the output of the superimposing coder 4.

The output of the index-value coder 5 is then applied as a row address pattern to the CMM 6, which already has data stored in it. As described above, the CMM outputs a column address pattern, the value of 10 each column representing the number of row intersections with that column, that are set to '1'.

The CMM output is fed to a threshold device 7, which provides an output in binary form, indicating the columns that meet the threshold value. The output is then fed to an MBI processor 8 (an example of which is 15 described below), which extracts all separator codes that match the output of the threshold device 7.

The extracted separators are then matched with their respective input data, to provide a result list 20. This result list 20 can then be subjected to a Back-Check operation (also described below), to match final results more 20 closely with the original query data.

An example of the above steps is summarised in Figure 5.

Figure 4 also shows a separator generator 9, which generates a separator code for each record that is entered into the CMM for storage, as mentioned above. The separator generator 9 generates separator codes 25 each having M bits, and each having the same number N ($N > 1$ and $N < M$) of

the M bits set to '1'. The N set bits of each separator code are randomly chosen (although no two separator codes are allowed to be identical), so that the various separator codes share a minimum number of set bits in common. (In some other applications, two or more separator codes may be
5 allowed to be identical.)

By choosing N to be greater than 1, M can be much less than it would be if N=1, and therefore the size of the respective CMM can be reduced, thus saving space. However, N could be 1 in some applications.

By allowing separator codes to have overlapping bits, a plurality of
10 possible matches to an input query may be obtained. The genuine match can be found from the possible matches by Back-Checking.

An alternative method to extract matches from a CMM query is "k-point" thresholding. Instead of selecting from the CMM output array of summed values those bits representing column values that are equal to or
15 greater than a predetermined numerical value ('2' in the above simple example), the k highest-value bits are selected, whatever those numerical values might be. In practice, a threshold is determined at the time of recall for each CMM that returns k or more bits set to one, to implement k-point thresholding. In practice, we usually find that k=N.

20 An improvement to assist finding an exact match of an input query is as follows. When binary codes representing tensored tuples are OR-ed (as in the above example for SPOTTER – Figure 2), the system can "lose" bits, for example:

Code 1: 0110000

25 Code 2: 0100010

Code 1 OR Code 2: 0110010

Note that the resultant OR-ed code (SIB, superimposed binding) has only 3 bits set, whereas the original codes have a total between them of 4 bits set. This can causes problems in the threshold stage of the CMM in
5 recall, in that too many possible matches may be listed.

The summed separator values that are output as an array in a subsequent recall or search process are a result of the above SIB being used to teach the CMM. The summed values are thresholded to obtain the separators of the possible matching data. For an exact match of the input
10 postal address to a stored postal address, the threshold can be set to the number of bits set in the input. But because of the "loss" of bits shown above, the system will give a lower threshold than "should" be given. In this case 3 instead of 4. This results in many more false hits from the memory.

A solution to this is to "multiple activate" in the teaching stage bits
15 which have two bits (or more) OR-ed on top of each other - i.e. those lines are counted more than once in the CMM teaching or access stage. In the recall or search stage, the threshold count then includes these multiple counts in the number of bits set in the input.

A number of further refinements can be incorporated into
20 embodiments of the invention.

Features may be incorporated into embodiments of the invention, to deal with skewed data. By skewed data is meant data in which certain items recur with a very high frequency, as compared to other items which occur very rarely. Storage of items of data with a high frequency of recurrence

can cause saturation of the CMM. This will be explained in more detail below.

As mentioned above, we have found that a particularly suitable technique for converting textual addresses into binary patterns for use with the CMM, was 3-tensoring or 'tripling' – that is, taking three characters of the input at a time, performing the tensoring operation for the binary tokens for each character, and converting the resulting 3-dimensional binary tensor product back into a 1-dimensional binary pattern. Other options are, for example, 2-tensoring (as described above) or 4-tensoring. However, it was found that with 2-tensoring, there were simply not enough combinations of characters to provide a large enough input to the CMM to avoid localised saturation. 4-tensoring would allow an even larger input to be generated, but then starts to suffer from intolerance to spelling mistakes, which is especially pronounced for smaller words which might be represented by only one or two quads. Tripling provides a reasonable compromise.

It was found that 37 individual characters could represent the textual addresses of the address data file. These were the 26 alphabetic characters, the 10 digits and the space. Thus, the binary tokens chosen were 37-bit binary patterns, each with 1 bit set, to give the maximum sparsity of code.

After tensoring, each triple would be represented by a binary pattern 50653 bits wide ($37 \times 37 \times 37$) with only 1 bit set. Thus each triple can be represented by a single number indicating the position within the 50653-bit wide binary pattern of the set bit. For example, the unique triple 'ROA' has the unique triple number of 39268, meaning that the binary pattern has bit 39268 set.

The input to the CMM is thus the activation of one CMM line for each triple in an address.

It was found that, for the address data of the address data file, the inputs to the CMM would be very badly skewed. The implications of skewed data are that particular lines of the CMM will be activated far more times than the average over all the CMM input lines. This means that some lines are used far more often than others. As data is only stored in the CMM on activated lines, this means that some lines are having to store far more data than others, rather than the data being spread evenly over the whole CMM.

5 This can lead to saturation of particular lines of the CMM (the ones which are activated far more often than others), and means that they not only fail to store more data, but also that they cannot reliably recall data already stored there.

10

One possible solution to this problem is simply to expand the CMM horizontally, making each line longer and longer, until there is sufficient capacity in each line to hold all of the information that will be stored there.

15 Obviously, as the CMM has to be rectangular, this means some lines will be far longer than necessary, in order to accommodate the frequently used lines. If the CMM is expanded so as to prevent saturation of the most commonly used input line, the majority of the CMM will be empty.

An analysis of the address data file shows that the most commonly used input triple occurs nearly 1.4 million times, while 91% of the input triples are used less than 10,000 times. This variation in the occurrence of the input triples means that some lines of the CMM are activated far more often than others, and leads to localised saturation of these lines.

One solution to this problem is to split the address data file into a number of smaller files according to some criteria and put each small file in its own smaller CMM.

As well as the input data from the address data file being far from uniformly distributed, another variation from the ideal profile of the CMM inputs is that the number of bits set in each input is not consistent from one address to another. This is because the addresses can be of varying length, and in 5 fact vary from 6 triples to 319 triples. In order to reduce this variation, it was decided that a good method for subdividing the address data file would be based on the address length. This means that for a particular file, and therefore a particular CMM, the number of input lines active would be more uniform. This process is referred to as 'banding'. The total number of 10 different address lengths is 313 (319-6), and by dividing this by the number of CMMs that it was planned to use, we arrive at a set of bands into which each address can be placed based on its length. For example, if we decide to use 3 CMMs, we get $313/3 \approx 104$. Therefore the bands will be 6-110, 111-215 and 216-319 triples. An address is placed in one of these bands according to how 15 many triples it contains.

In the limit, we could decide to use 313 CMMs, which would mean that each separate address length would have its own CMM. However, analysis shows that, even now, the worst case (which happens with addresses of length 12 triples) requires a threshold of 12860. This means that the 20 separator would have to be 25720 bits wide, and one address would still have only a single triple stored in the CMM. Successful recall would rely upon that triple being included in the query.

A better method for allocating multiple CMMs to the problem was required. Instead of trying to split the file up into well defined blocks, it was 25 noted that a more efficient method would be to add addresses to the CMM until a particular line became too saturated to take new data. Address data is then allocated to the CMMs in the following manner.

1. Start with one CMM.
2. Determine whether the first address to be stored would access any lines of the first CMM which are already saturated.
3. If none of the accessed lines is saturated, store the address in the first CMM and continue from Step 5.
4. If at least one of the lines is saturated, store the address in a new, empty CMM.
5. Read the next address to be stored and determine whether it would access any lines of the first CMM which are already saturated.
- 10 6. If none of the accessed lines is saturated, store the address in the first CMM and continue from Step 5.
7. If at least one of the lines is saturated in the first CMM, continue checking all the other CMMs to find one where none of the accessed lines is saturated.
- 15 8. If a CMM is found, for which none of the accessed lines is saturated, store the address at this CMM and continue from Step 5.
9. If no CMM is found, for which none of the accessed lines is saturated, store the address in a new, empty CMM, and continue from Step 5.
- 20

This process will ensure that no particular CMM line exceeds a chosen level of saturation, and that each address is stored in the first available CMM without exceeding this saturation level.

The problem of skewed data may not be overcome completely by simply splitting up the file, and it was desirable not to exclude data simply because it was too common. Taking these two requirements resulted in very wide CMMs where only a few lines were very heavily used. While this 5 approach should in theory perform well, its drawback is the amount of memory required to hold the CMM. Given the skewedness of the input distribution, it is clear that the majority of the CMM would be filled with zeros. This means that the majority of the CMM is not actually holding any information and it follows from this that it ought to be possible to compress the CMM, reducing the 10 amount of space wasted while not affecting the parts of the CMM which actually hold useful data. A technique for achieving this was developed – by dynamically altering the implementation of the CMM on a per-line basis, according to how much data the line has to store.

The conventional method for implementing a CMM line has been to 15 hold the binary pattern for that line as an array of words, as shown in Figure 6a. This shows a single 24-bit CMM line which has a single 2-bit-set separator in it, and its representation as an array of three 8-bit words. However, an alternative representation is shown in Figure 6b. This new method of implementing a CMM line is to store a list of the positions of the bits set to 1. 20 These are the only significant parts of the CMM, and this removes the need to actually store all of the zeros as well. In this example, the CMM is 24 bits wide. This means the highest bit position is 23 (starting from 0), and the number 23 can be represented in 5 binary bits. So in this case, a 24-bit CMM line with a single 2-bit separator stored in it can be represented using only 10 25 bits – the two 5-bit numbers indicating the positions of the set bits in the CMM line.

The implications are not immediately obvious for such a small CMM line, but the table in Figure 7 shows how many 2-bit-set separators can be stored using this new method for larger CMMs. This table assumes a machine word size of 16 bits. If the number of separators to be stored in a

5 line exceeds the usage threshold, then it becomes more memory-efficient to store a real binary representation of the CMM line. The 3rd and 4th columns show the total memory requirement for a 20000 input line CMM, where 10% of the lines are used 400 times, and the other 90% of the lines are used only 20 times each.

10 There are a few points which should be noted about the figures presented in Figure 7.

1. The method of storing the CMM line, whether conventional binary or compressed, makes no difference to the operation of the CMM as a whole. It is quite possible to mix ordinary and compressed lines within one CMM. The only difference is that a conventional line must be scanned for set bits during the recall phase in order to accumulate the summed values, whereas a compressed line already has the bits positions implicitly stored.
2. The values of 10%, 90%, and their corresponding usage counts of 400 and 20 were chosen to illustrate the point that when most of the inputs fall below the usage threshold value, a considerable saving in memory can be obtained, especially as the CMM is made larger.
3. The amount of memory required by the compressed CMM lines in this example is constant, no matter what the width of the CMM. This is the main reason that the 10000-bit wide CMM

shows such a large memory saving. In effect, only 10% of the CMM has grown to 10000 bits wide, while the other 90% is still using the same amount of memory as it was for a 1000 bit wide CMM.

5 4. In a real application, the number of times each input line is activated is unlikely to fall into 2 groups as in the example here. In practice, the usage count for each input line can be seen to climb very slowly until it suddenly grows very rapidly indeed. The CMM width should be chosen to be as large as is practical, 10 so that most of the lines will be stored compressed, and only the very commonly used lines will be stored as binary CMM lines.

15 By running an analysis of the address database using these new techniques, the data shown in Figure 8 was obtained. It shows the total amount of memory required to store the full address database using various widths of CMM.

It can be seen from this table that there is a minimum memory requirement somewhere around the 15000 bit size. In fact, going down to 1000-bit intervals, this point is reached at 14000 bits. What this implies is that the advantages of storing compressed CMM lines is maximised - using any 20 wider CMMs results in the extra overhead of holding very wide real binary lines outweighing the savings which can be made using compressed lines. From this value, we can work out how many separators can be stored in a CMM line before it becomes more economical to store them as real binary CMM lines. 14000 bits gives us 1750 bytes of storage. Each bit position in a 25 separator would require 2 bytes, and each separator has 2 bits set. So dividing this by 4 gives us a value of 437 separators. This means that if a particular line of a particular CMM is used 437 or fewer times, it would be

more memory efficient to store it as a compressed line. Further analysis shows that for this particular application, 198 memories are required, ranging in size from 12264x14000 bits down to 3585x14000 bits. On average, each memory has 3.7% real binary lines and 96.3% compressed lines. Of these,

- 5 nearly 20% of the inputs to each memory are used only once during training. The total memory requirement for this configuration is 291.2Mb. As a comparison, if each CMM line was stored as a "real" binary line, the total memory requirement would be 2.4Gb. The compression technique has therefore achieved over an eight-fold memory saving in this particular case.
- 10 As was implied earlier, the wider the CMM, the greater the relative memory saving that can be obtained using compressed CMM lines. The table in Figure 9 shows the relative savings for the CMMs detailed previously in Figure 8.

Further pre-processing of the text may be carried out, to help in

- 15 reducing the incidence of saturation of CMM lines. In the example of the address data file, this may be achieved by removing hyphen and number characters from an input data string, except for those characters forming part of the postcode. That is, the string is processed to remove unwanted characters and generate a set of tokens (i.e. sub-strings deemed to be valid
- 20 inputs for the purpose of generating CMM input codes). The potential loss of information can then be made up for by more accurate matching which takes place in a Back-Check function.

An example of such a tokenisation process is as follows:

- 25 String to tokenise:
"68 45636 1-23 31-49 SANDRINGHAM ROAD SOUTHAMPTON SO18 1JL"
Found 9 tokens

	Token 0: "68"	Remove unacceptable string from input
	Token 1: "45636"	Remove unacceptable string from input
	Token 2: "1-23"	Remove unacceptable string from input
	Token 3: "31-49"	Remove unacceptable string from input
5	Token 4: "SANDRINGHAM"	Accepted
	Token 5: "ROAD"	Accepted
	Token 6: "SOUTHAMPTON"	Accepted
	Token 7: "SO18"	Accepted
	Token 8: "1JL"	Accepted

10 As outlined above, a unique binary code is then assigned to each unique input character. It is useful to provide access to the mapping between characters and binary codes so that applications can "look-up" the character/string for a given code, and vice-versa. The actual form of the code is determined by three parameters. One parameter specifies the width of the
15 bit field (in bits) for all codes to be used in a particular CMM. A second parameter specifies the number of bits which are to be set to logical '1' in the code, which is normally a small, fixed number. A third parameter may be provided which specifies the minimum permitted Hamming distance between any pair of codes used in a particular CMM. This provides some control over
20 the amount of "overlap" between codes used and helps to minimise spurious outputs during subsequent recall. The codes may then be processed and stored in the CMM, as described above.

25 Embodiments of the invention may use Middle-Bit-Indexing (MBI) extraction to extract separators from a superimposed separator (conveniently referred to as an SIS) as may be produced as the binary output from the CMM, as described above. An example of an MBI extractor is illustrated in Figure 10.

The MBI extractor illustrated in Figure 10 is very efficient in maintaining mappings between separators and the records they represent and, most importantly, to process an SIS by extracting any individual separators that are present, and retrieving the corresponding records.

5 The MBI extractor uses the middle bits of the SIS to determine which buckets of a separator database to access during search. Each separator is stored in a bucket corresponding to the location of the middle bit of the bits that are set (to '1') in that separator.

Basically, the "middle bits" of a separator are used to identify the
10 bucket where that separator would be stored in the separator database if the separator in question existed. The problem is that, during a recall operation, an SIS is obtained which contains a number of possible separator codes, of which some may not exist in the system. (Recall that a separator is only brought into existence if it is created when entering a record into the CMM.)
15 MBI uses an index into the separator database based on the position of the middle bit of existing separator codes, to find the genuine separators in an SIS and hence the records represented by those genuine separators.

Figure 10 shows the type of data structure used in MBI extractors.
Here the buckets are enumerated such that the ordinal number of the bucket
20 corresponds with the position of the middle set bit for those separators stored within that bucket (counting from left to right) in the separator database. Buckets 0 and 6 are not shown, since they are not used in this MBI-based example. This is simply because the separators are seven bits wide and bits 0 and 6 could not be "middle bits".

25 Consider bucket 5, for example. This bucket has just one entry for the separator 1000011, which has the 5th bit set. If the separator 0010011

existed in this system, it too would be stored in bucket 5. Note that the separator is not stored explicitly in this extractor (that is, as the full binary number such as 1100001), but rather as an array of integers representing the set bit positions in the separator (that is, the corresponding shorter array 5 such as [0,6]). Since the bucket number is the same as the bit position of the middle bit for all separators in a particular bucket, the integer representing the fact that the middle bit is set is omitted from the array of bits set to avoid redundancy.

During a recall operation, the SIS generated from the CMM is 10 inspected to identify each bit position which represents a potential middle bit in a separator. A group of one or more set bits on the extreme left and right are discounted immediately, because they cannot assume a middle position. For example, if the separators are fixed to have always 5 bits set, then the 2 bits on the extreme left and right of the SIS can never be middle bits.

15 Once the extreme bit positions have been discounted, every other set bit is a potential middle bit and the corresponding buckets must be checked. One implementation of this uses AND separator checking. This uses a bitwise logical AND between each separator stored in each selected bucket and the SIS. If a stored separator is unchanged by the AND operation, the 20 SIS must contain that separator. The number (identifier) of each found separator is added to a list so that the records represented by the separators can be subsequently recovered.

The result of a search using the CMM are sometimes imperfect in the sense that some of the separators appearing in the SIS are spurious and 25 merely a side-effect of the way CMM storage operates. This means that some records may be erroneously indicated as matching the input. A Back-

Checker is a device which aims to verify whether each result record is a true possible match with the input or whether it is a spurious result (non-match).

One implementation of a Back-Checker operates by counting the number of words in every result record which match a word in the input query. In addition, the score is modified by including the result of comparing the soundex code of words in the address. A soundex code represents the sound of a word such that similar sounding words are meant to have the same soundex code. Increasing the score for a matching soundex code is intended to improve tolerance to minor spelling errors in the query. The resulting count is used to rank the results according to how well they match the input.

A particular advantage of the illustrated data processors is that a significant amount of data compression takes place in the storage of indexes in a CMM. This is a consequence of using multiple set bits in the separator code (i.e. more than 1 bit is set in each separator code), and the columnwise summation of bits from selected rows. The former means that separator codes may be overlapped during the training phase (i.e. more than one separator code can have the same bit set), and the latter means that any aliasing introduced by this overlapping can be resolved during recall. It may be noted that summation results in a kind of voting system, where selected rows cast a vote for a particular separator/record at the places where a bit is set. This provides a much more selective system than, for example, many previous methods which have (in effect) used a logical AND to combine bits from selected rows. The number of overlapping codes increases with the number of separators stored, and eventually errors may occur during recall in the form of additional records being selected as well as the genuine matches. This is also seen as a feature of the system, where one can trade-off

increasingly compressed storage of indexes against a small (but increasing with compression) probability of erroneous records together with the desired records.

Those skilled in the art will readily appreciate that rows and columns 5 of a matrix can readily be interchanged. For example, binary patterns and separators that are illustrated as being entered respectively as rows and columns could equally be entered respectively as columns and rows, provided that all data entry and recall is consistent in the convention chosen.

Although the illustrated embodiments of the invention have been 10 described by way of example as for address matching in the address data file, it is to be appreciated that other embodiments of the invention may be used to store and recall data of any other type.

In this specification, the verb "comprise" has its normal dictionary meaning, to denote non-exclusive inclusion. That is, use of the word 15 "comprise" (or any of its derivatives) to include one feature or more, does not exclude the possibility of also including further features.

The reader's attention is directed to all papers and documents which are filed concurrently with or previous to this specification in connection with this application and which are open to public inspection with this 20 specification, and the contents of all such papers and documents are incorporated herein by reference.

All of the features disclosed in this specification (including any accompanying claims, abstract and drawings), and/or all of the steps of any method or process so disclosed, may be combined in any combination,

except combinations where at least some of such features and/or steps are mutually exclusive.

Each feature disclosed in this specification (including any accompanying claims, abstract and drawings), may be replaced by 5 alternative features serving the same, equivalent or similar purpose, unless expressly stated otherwise. Thus, unless expressly stated otherwise, each feature disclosed is one example only of a generic series of equivalent or similar features.

The invention is not restricted to the details of the foregoing 10 embodiment(s). The invention extends to any novel one, or any novel combination, of the features disclosed in this specification (including any accompanying claims, abstract and drawings), or to any novel one, or any novel combination, of the steps of any method or process so disclosed.

CLAIMS

1. A data processor comprising:
 - a) a correlation matrix memory arranged to store data;
 - 5 b) input means arranged to receive sets of input data to be stored in the correlation matrix memory;
 - c) a sampler arranged to derive, from each set of input data, a respective set of tuples;
 - d) a coder arranged to code each of the tuples;
 - 10 e) a combiner arranged to combine the coded tuples for a respective set of input data;
 - f) a separator generator arranged to generate for each set of input data a respective, associated, unique separator;
 - 15 g) storage means arranged to store the association of each separator with its respective set of input data; and
 - h) addressing means arranged to applying to the correlation matrix memory, for each set of input data, the respective combined coded tuples as a row address and the respective unique separator as a column address, or vice-versa.
- 20 2. A data processor according to claim 1, wherein the combined coded tuples for each set of input data are in the form of a binary coded word; the data processor further comprises a translator arranged to translate each such binary coded word into a translated word

comprising index values representing which bits of the binary coded word are set; and said addressing means is arranged to apply the translated word to the correlation matrix memory.

3. A data processor according to claim 1 or 2, wherein said separator generator is arranged to generate separators in a random manner.
5. A data processor according to claim 1, 2 or 3, wherein said separator generator is arranged to generate separators which are M bits wide and having N bits set, where $N > 1$ or $N = 1$, and where $N < M$.
10. A data processor according to any of the preceding claims, wherein, for each said set of tuples, each tuple comprises three successive elements of a respective set of input data, and each successive tuple is offset by one such element from the preceding tuple.
15. 7. A data processor according to any of the preceding claims, wherein said combiner is arranged to combine the coded tuples for a respective set of input data, by superimposition.
20. 8. A data processor according to any of the preceding claims, wherein at least some of the rows (or columns) of the correlation matrix memory are represented by binary words, each of which represents the positions of each bit in the respective row (or column) which is set.
25. 9. A data processor according to any of the preceding claims, wherein said correlation matrix memory comprises a plurality of sub-

correlation matrix memories; said addressing means is arranged to access a first one of said sub-correlation matrix memories and apply the combined coded tuples of a respective set of input data to that sub-correlation matrix memory unless a respective row (or column) of that sub-correlation matrix memory will become saturated by application of those tuples; and in the event of such prospective saturation, access successive ones of the sub-correlation matrix memories until those tuples can be applied to a respective one of the sub-correlation matrix memories without such saturation.

5 10. A data processor according to any of the preceding claims, arranged to receive sets of query data to be matched with sets of input data stored in the correlation matrix memory, and to derive, for each set of query data, a respective set of coded tuples analogous to those derived for the original input data, and to apply to the correlation matrix memory, for each set of query data, the respective combined coded tuples as a row (or column) address: the data processor further comprising:

15 a) output means for outputting a raw superimposed separator which represents, for a respective set of query data, the number of rows (or columns) having a bit set by the applied combined coded tuples in each column (or row) represented by the raw superimposed separator;

20 b) threshold means arranged to convert the raw superimposed separator into a binary superimposed separator; and

25 c) an extractor arranged to extract individual separators from the binary superimposed separator.

11. A data processor according to claim 10, wherein said thresholding means sets an absolute threshold value, and provides said binary superimposed separator as a word in which bits represent respective columns (or rows) of the correlation matrix memory, and each of those bits is set if the number of rows (or columns) having a bit set by the applied combined coded tuples in the respective column (or row) equals or is greater than said absolute threshold value.
12. A data processor according to claim 10, wherein said thresholding means determines a value k, and provides said binary superimposed separator as a word in which bits represent respective columns (or rows) of the correlation matrix memory, and are set for the k respective columns (or rows) having the highest number of rows (or columns) which have a bit set by the applied combined coded tuples in the correlation matrix memory.
13. A data processor according to any of claims 10 to 12, further comprising back-checking means arranged to compare sets of recalled data, identified by respective separators extracted by said extractor, with original query data, in order to identify the set or sets of recalled data which matches best the original query data.
14. A data processor according to any of the preceding claims, arranged to process sets of input data and query data in the form of postal addresses.
15. A data processor substantially as hereinbefore described with reference to the accompanying drawings.

16. A data processor substantially as hereinbefore described with reference to Figure 4 of the accompanying drawings.
17. A method of processing data, comprising the steps of:
 - a) receiving sets of input data to be stored in a correlation matrix memory;
 - b) deriving, from each set of input data, a respective set of tuples;
 - c) coding each of the tuples;
 - d) combining the coded tuples for a respective set of input data;
 - e) generating for each set of input data a respective, associated, unique separator;
 - f) storing the association of each separator with its respective set of input data; and
 - g) applying to the correlation matrix memory, for each set of input data, the respective combined coded tuples as a row address and the respective unique separator as a column address, or vice-versa.
18. A method of processing data comprising the steps of:
 - a) receiving sets of query data to be matched with sets of input data stored in a correlation matrix memory by a method according to claim 17;
 - b) deriving, for each set of query data, a respective set of coded tuples analogous to those derived for the original input data;

- c) applying to the correlation matrix memory, for each set of query data, the respective combined coded tuples as a row (or column) address;
- d) outputting a raw superimposed separator which represents, for a respective set of query data, the number of rows (or columns) having a bit set by the applied combined coded tuples in each column (or row) represented by the raw superimposed separator;
- e) converting the raw superimposed separator into a binary superimposed separator;
- f) extracting one or more individual separator from the binary superimposed separator; and
- g) identifying the or each respective original set of input data from association with the or each individual separator extracted from the binary superimposed separator.

19. A method according to claim 17 or 18 and carried out by a data processor according to any of the claims 1 to 16.

20. A method according to claim 17, 18 or 19 and incorporating any of the features disclosed in the accompanying specification and/or drawings.

21. A method of processing data substantially as hereinbefore described with reference to the accompanying drawings.

22. A method of processing data substantially as hereinbefore described with reference to Figure 4 of the accompanying drawings.

THIS PAGE BLANK (USPTO)

1 / 5

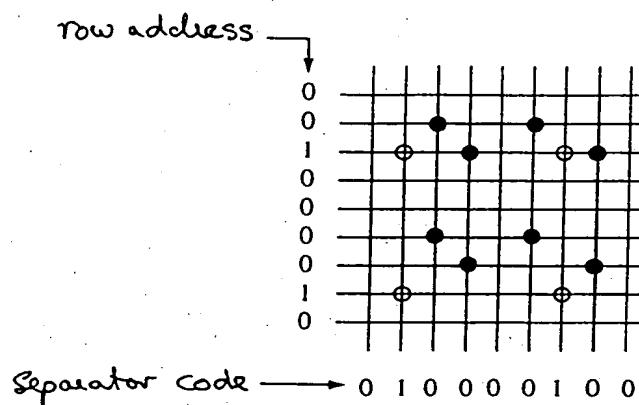
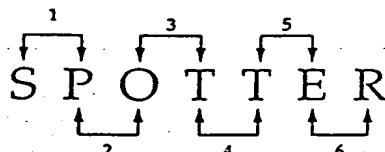


FIG. 1

(a) Character Tokens:	S	100000
	P	010000
	O	001000
	T	000100
	E	000010
	R	000001

(b) Tensoring Order:



(c) Tensor Pattern 1:

(d) Pattern for 'Spotter': 000000100000010000001100000100000010

(e) Pattern for 'Spoter': 000000100000010000001000000100000010

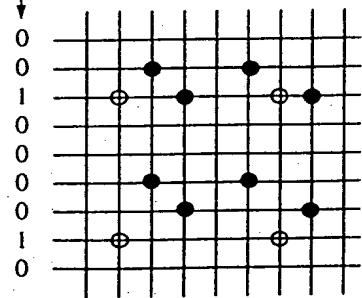
(f) Pattern for 'Protest': 00001000000000001101000000100010000

FIG. 2

THIS PAGE BLANK (USPTO)

2/5

query data



Sum output: 0 2 0 1 0 0 2 1 0

Thresholded sum → 0 1 0 0 0 0 1 0 0

FIG. 3

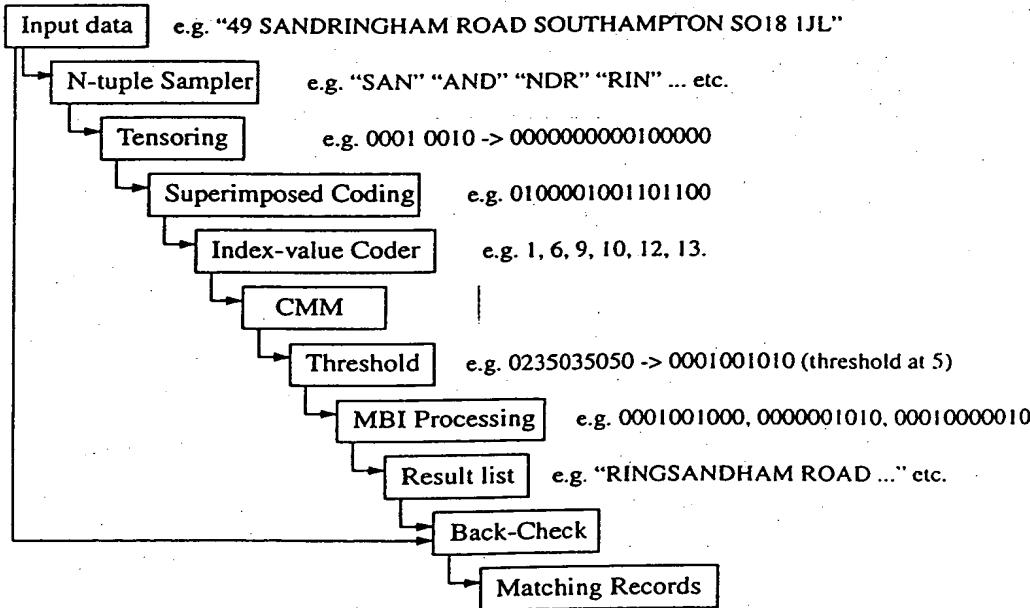
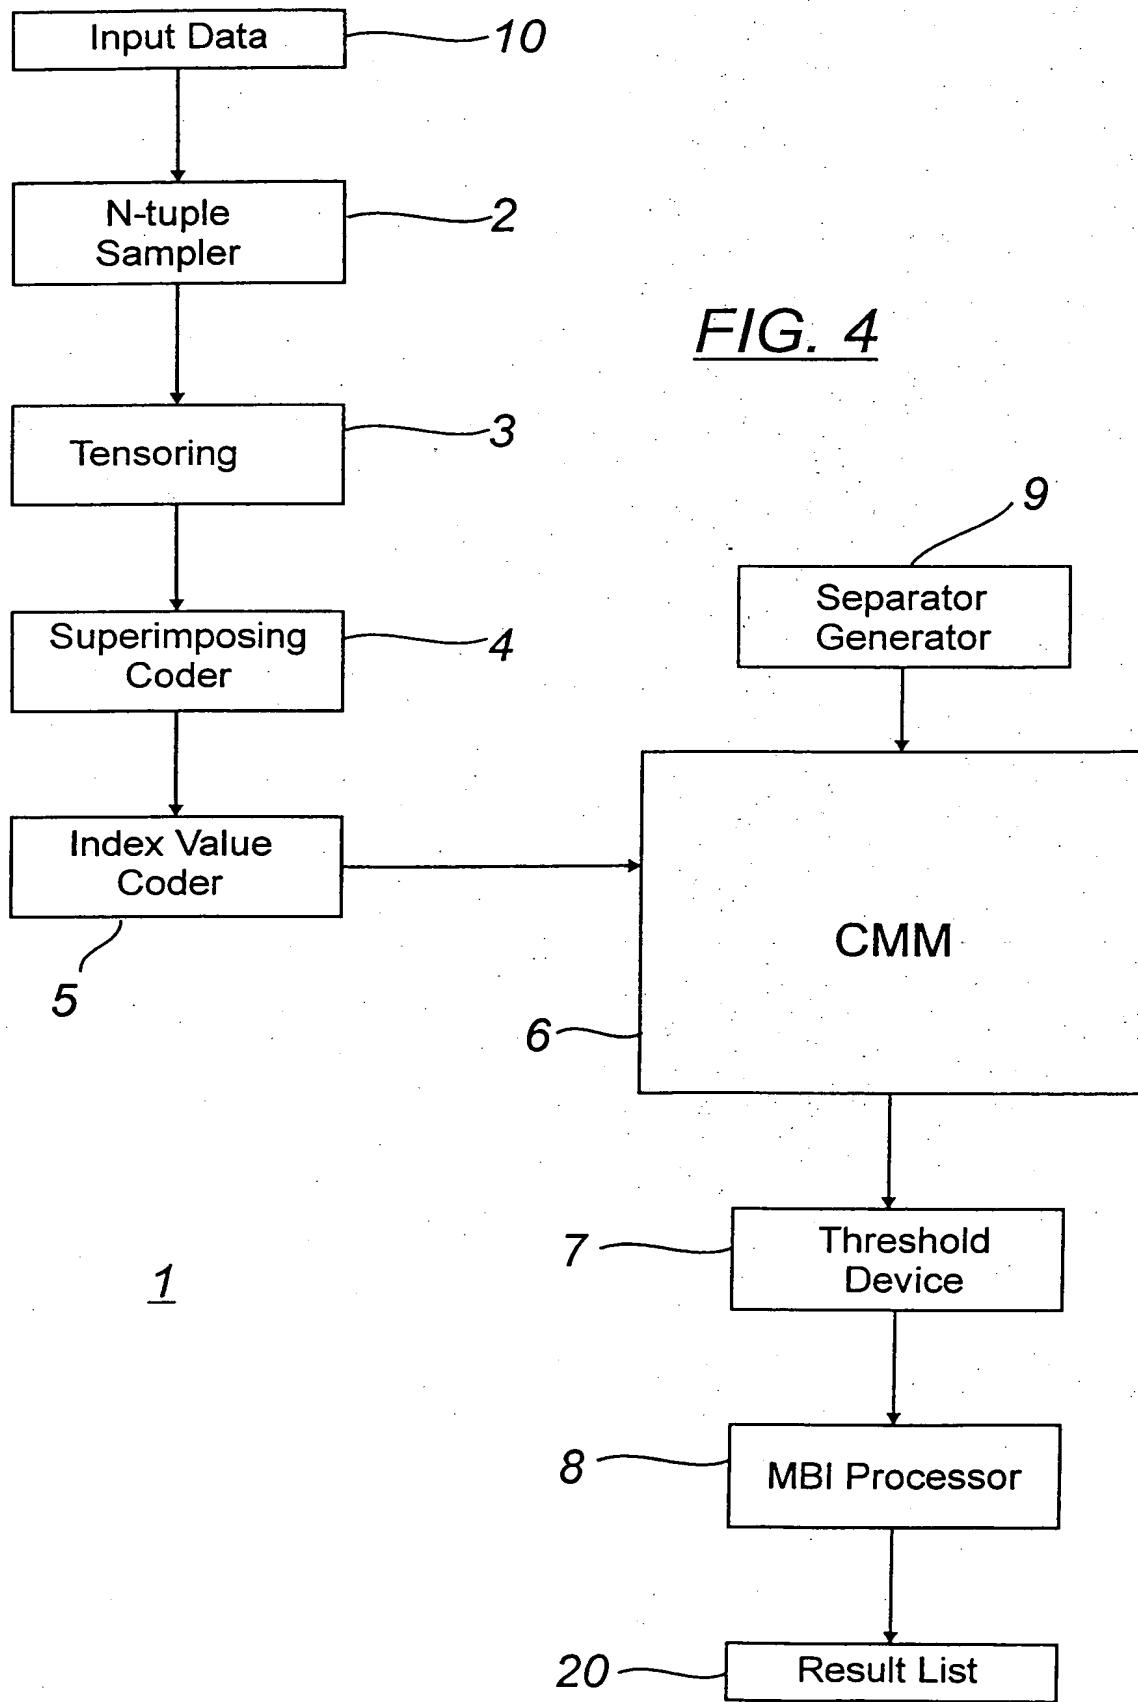


FIG. 5

THIS PAGE BLANK (USPTO)



THIS PAGE BLANK (USPTO)

4/5

CMM Line	00010000	00000000	10000000
Word Array	00010000	00000000	10000000

FIG. 6a

CMM Line	00010000	00000000	10000000
Bit Set Position	3	16	

FIG. 6b

CMM Width (Bits)	Usage Threshold (Separators)	Pure CMM (Mb)	Compressed CMM (Mb)
1000	31	2.38	1.61
2000	62	4.77	1.85
5000	156	11.92	2.57
10000	312	23.84	3.76

FIG. 7

Width of CMM (bits)	Memory required (Mb)	Total Number of Memories
5000	305.45	553
10000	292.31	277
15000	291.38	185
20000	292.48	139
25000	294.13	111
30000	295.63	93
35000	297.21	79
40000	298.88	70

FIG. 8

THIS PAGE BLANK (USPTO)

5/5

Width of CMM (bits)	Memory Requirement Using Compression (Mb)	Memory Requirement Using Binary CMM (Mb)	Relative Memory Saving
5000	305.45	1770.17	5.8
10000	292.31	2229.29	7.6
15000	291.38	2498.57	8.6
20000	292.48	2689.14	9.2
25000	294.13	2832.08	9.6
30000	295.63	2945.26	10
35000	297.21	3036.78	10.2
40000	298.88	3131.58	10.5

FIG. 9

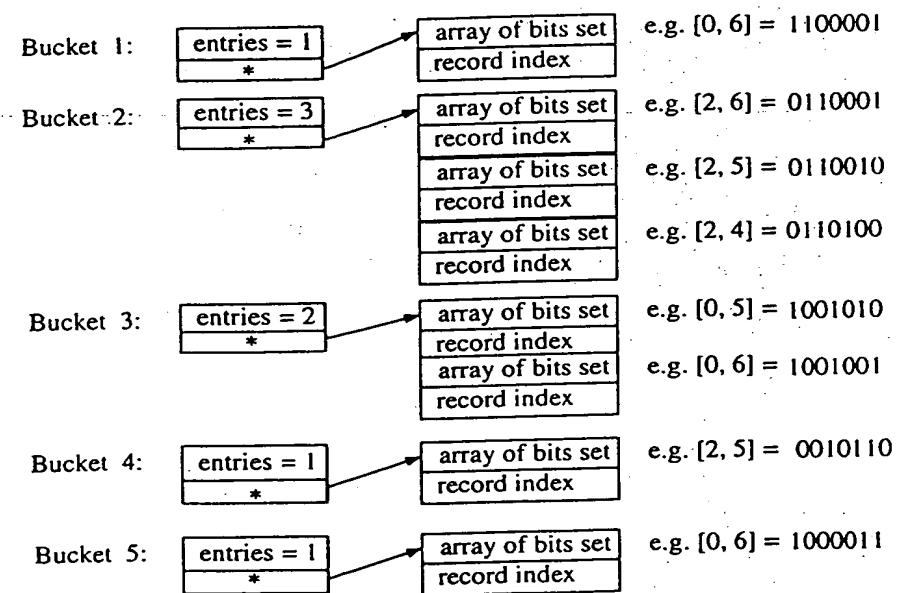


FIG. 10

PCT 19B00102303

Stanley Gordon
25-7-00

THIS PAGE BLANK (USPTO)